Claims

What is claimed is:

1. A method of serving data to a plurality of clients in a client-server environment, comprising the steps of:

providing a plurality of versions of data in which at least two versions have different overheads associated therewith;

assigning individual clients to one of a plurality of quality-of-service classes; and

satisfying requests so that a client belonging to a high quality-of-service class is given preferential access to data versions which require higher overheads to serve.

2. The method of claim 1, wherein the overhead to serve a version is correlated with a quality of the version.

3. The method of claim 2, wherein the multiple versions comprise images of different resolutions and clients belonging to a high quality-of-service class are given preferential access to higher resolution images.

4. The method of claim 2, wherein the quality of a version is correlated with a processing time required to create the version.

5. The method of claim 1, wherein the overhead to serve a version is correlated with how current the version is.

6. The method of claim 1, further comprising the step of:

in response to a system load exceeding a threshold, satisfying a higher percentage of requests from clients belonging to a lower quality-of-service class with a version requiring lower overhead to serve.

7. The method of claim 1, wherein the server comprises multiple nodes and different nodes provide data versions requiring different overheads to serve.

8. The method of claim 1, further comprising the step of implementing a quality-of-service policy that specifies at least one of content quality and latency.

9. The method of claim 8, wherein one or more clients belonging to a premium service class are served with high content quality and low latency.

10. The method of claim 8, wherein one or more clients belonging to a medium service class are served with one of high content quality and low latency.

11. The method of claim 8, wherein one or more clients belonging to a best-effort service class are served with unspecified content quality and latency.

12. The method of claim 1, wherein a client request is routed using at least one of an identity of the client, a quality of content, a load on at least one server, a data distribution on at least one server, and a capacity of at least one server.

13 The method of claim 1, wherein a client is assigned to a quality-of-service class by program logic that is externalized from the server.

14 The method of claim 13, wherein the externalized program logic comprises a set of business rules that can be modified by nonexperts in information technology.

15. The method of claim 1, further comprising the step of satisfying requests using a policy determined by program logic that is externalized from the server.

16. The method of claim 15, wherein the externalized program logic comprises a set of business rules that can be modified by nonexperts in information technology.

17. Apparatus for serving data to a plurality of clients in a client-server environment, comprising:

a memory, and

at least one processor coupled to the memory and operative to: (i) provide a plurality of versions of data in which at least two versions have different overheads associated therewith; (ii) assign individual clients to one of a plurality of quality-of-service classes; and (iii) satisfy requests so that a client belonging to a high quality-of-service class is given preferential access to data versions which require higher overheads to serve.

18. The apparatus of claim 17, wherein the overhead to serve a version is correlated with a quality of the version.

19. The apparatus of claim 18, wherein the multiple versions comprise images of different resolutions and clients belonging to a high quality-of-service class are given preferential access to higher resolution images.

20. The apparatus of claim 18, wherein the quality of a version is correlated with a processing time required to create the version.

21. The apparatus of claim 17, wherein the overhead to serve a version is correlated with how current the version is.

22. The apparatus of claim 17, wherein the at least one processor is further operative to, in response to a system load exceeding a threshold, satisfy a higher percentage of requests from clients belonging to a lower quality-of-service class with a version requiring lower overhead to serve.

23. The apparatus of claim 17, wherein the at least one processor comprises multiple nodes and different nodes provide data versions requiring different overheads to serve.

24. The apparatus of claim 17, wherein the at least one processor is further operative to implement a quality-of-service policy that specifies at least one of content quality and latency.

25. The apparatus of claim 24, wherein one or more clients belonging to a premium service class are served with high content quality and low latency.

26. The apparatus of claim 24, wherein one or more clients belonging to a medium service class are served with one of high content quality and low latency.

27. The apparatus of claim 24, wherein one or more clients belonging to a best-effort service class are served with unspecified content quality and latency.

28. The apparatus of claim 17, wherein a client request is routed using at least one of an identity of the client, a quality of content, a load on at least one server, a data distribution on at least one server, and a capacity of at least one server.

29. An article of manufacture for use in serving data to a plurality of clients in a client-server environment, comprising a machine readable medium containing one or more programs which when executed implement the steps of:

providing a plurality of versions of data in which at least two versions have different overheads associated therewith;

assigning individual clients to one of a plurality of quality-of-service classes; and

satisfying requests so that a client belonging to a high quality-of-service class is given preferential access to data versions which require higher overheads to serve.

30. A system, comprising:

a plurality of clients, each client belonging to a quality-of-service class;

a load balancer for sending requests from clients to at least one back-end server; and

at least one back-end server for providing a plurality of versions of different objects in which at least two versions of an object have different overheads associated therewith.

31. A method of providing a data serving service, comprising the step of:

a service provider: (i) providing a plurality of versions of data in which at least two versions have different overheads associated therewith; (ii) assigning individual clients to one of a plurality of quality-of-service classes; and (iii) satisfying requests so that a client belonging to a high quality-of-service class is given preferential access to data versions which require higher overheads to serve.

32. The method of claim 31, wherein the data serving service comprises a quality-of-service policy specification.

33. The method of claim 32, wherein the quality-of-service policy specification comprises:

a plurality of subscriptions, each subscription being specified by content quality and service latency, wherein a limited premium service subscription is served with high content quality in low service latency, a medium service subscription is served with a high content quality or a low service latency, and an unlimited best-effort service subscription is served with unspecified content quality and latency.

34. The method of claim 31, wherein the service provider modifies data content and how the data content is served to clients in response to one or more changing conditions.

35. The method of claim 34, wherein one or more changing conditions comprises a source of a bottleneck.

36. The method of claim 31, wherein the step of assigning individual clients to one of a plurality of quality-of-service classes is based on a client payment.

37. A method of serving data to a plurality of clients, comprising the steps of:

establishing at least two quality-of-service classes; and

satisfying requests so that a client belonging to one quality-of-service class is served with a data version having one overhead associated therewith, while a client belonging to another quality-of-service class is served with a data version having another overhead associated therewith.

5